



US005689645A

United States Patent [19]

[11] Patent Number: 5,689,645

Schettler et al.

[45] Date of Patent: Nov. 18, 1997

[54] PERSISTENCE SPECIFICATION SYSTEM  
AND METHOD FOR PRODUCING  
PERSISTENT AND TRANSIENT SUBMAPS IN  
A MANAGEMENT STATION FOR A DATA  
COMMUNICATION NETWORK

5,463,628 10/1995 Sorensen ..... 370/110.1  
5,473,608 12/1995 Gagne et al. .... 370/85.13  
5,488,715 1/1996 Wainwright ..... 395/182.02  
5,491,796 2/1996 Wanderer et al. .... 395/200.09  
5,509,123 4/1996 Dobbins et al. .... 395/200.15  
5,539,886 7/1996 Adred et al. .... 395/200.04

[75] Inventors: Robert Dwight Schettler; William  
Girard McCollom; David M.  
Haimson, all of Fort Collins, Colo.

Primary Examiner—Alpesh M. Shah

[73] Assignee: Hewlett-Packard Co., Palo Alto, Calif.

[21] Appl. No.: 348,024

[22] Filed: Dec. 1, 1994

[51] Int. Cl.<sup>6</sup> ..... G06F 13/00

[52] U.S. Cl. .... 395/200.1; 395/100

[58] Field of Search ..... 395/200.01, 200.02,  
395/200.1, 200.2, 827, 878, 402, 444, 480,  
500, 550, 600, 800, 100, 118, 133

[56] References Cited


## U.S. PATENT DOCUMENTS

5,049,873 9/1991 Robins et al. .... 340/825.06  
5,063,523 11/1991 Vrenjak ..... 364/514 C  
5,123,098 6/1992 Gunning et al. .... 395/402  
5,136,690 8/1992 Becker et al. .... 395/161  
5,185,860 2/1993 Wu ..... 395/200  
5,276,789 1/1994 Besaw et al. .... 395/140  
5,414,812 5/1995 Filip et al. .... 395/600  
5,453,980 9/1995 Van Engelshoven ..... 370/60.1  
5,455,952 10/1995 Gjovaag ..... 395/700

## [57] ABSTRACT

Discovery/layout software configures a general purpose computer system to act as a management station using an industry standard SNMP protocol. The discovery/layout software has a discovery mechanism and a layout mechanism which, in combination, permit the discovery/layout software to provide various submaps to a display for illustrating network topology, which includes devices and device interconnections of the network. The submaps correspond to various hierarchical views of the network. Significantly, a persistence specification mechanism is provided in the discovery/layout software for specifying a submap as either transient (generated upon demand) or persistent (exists whether demanded or not). An integrating application as well as the user can identify a submap as persistent. This feature enables better interfacing of the integrating application with the station, thereby providing more information to the user. This feature further minimizes memory requirements as well as requisite processing time due to the elimination of unnecessary submaps and the elimination of processing of topology changes relative to the unnecessary submaps.

13 Claims, 18 Drawing Sheets







FIG 1

**FIG 2**




FIG 3

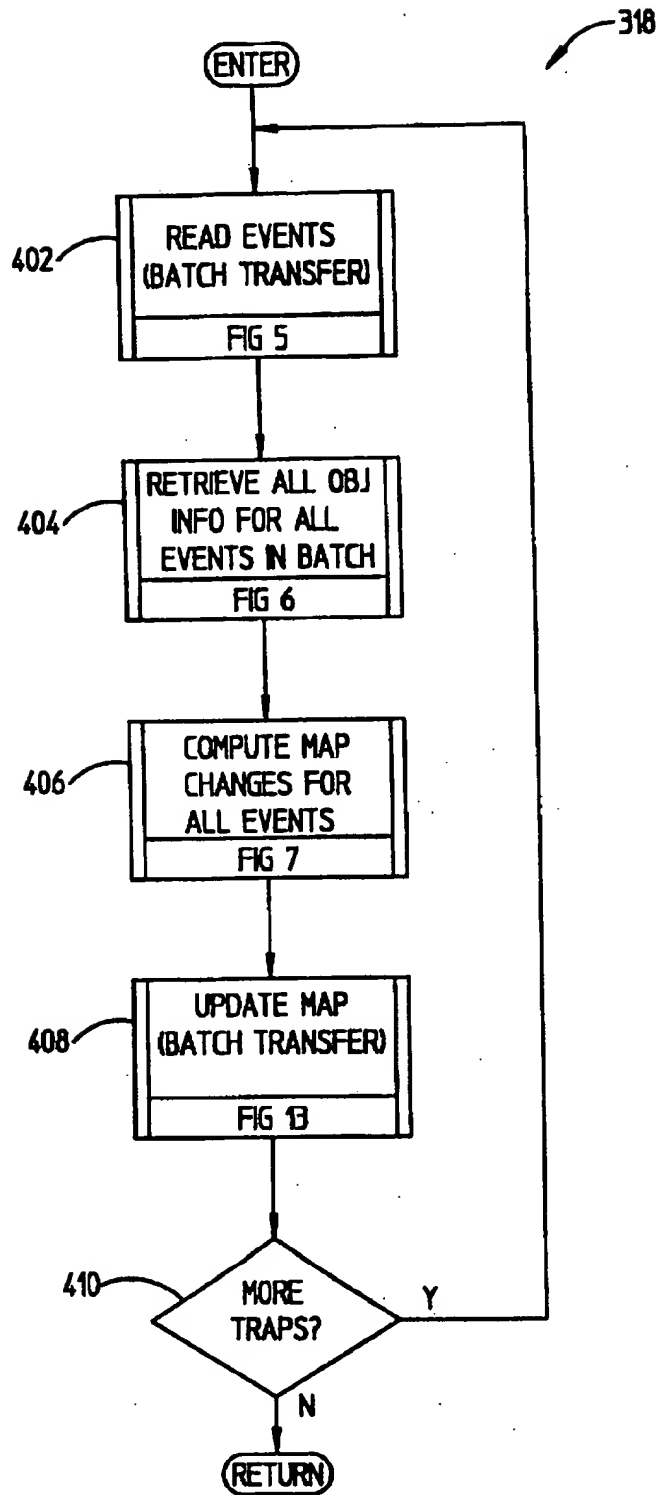
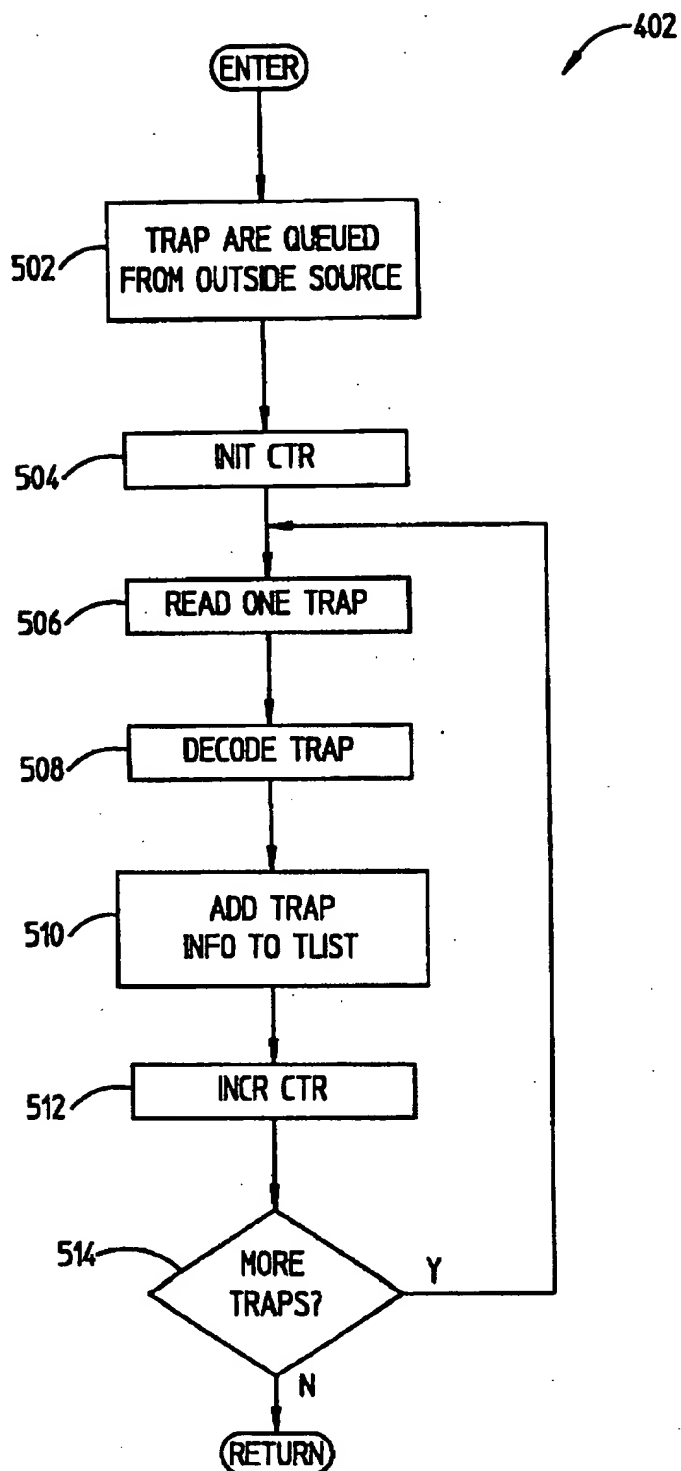
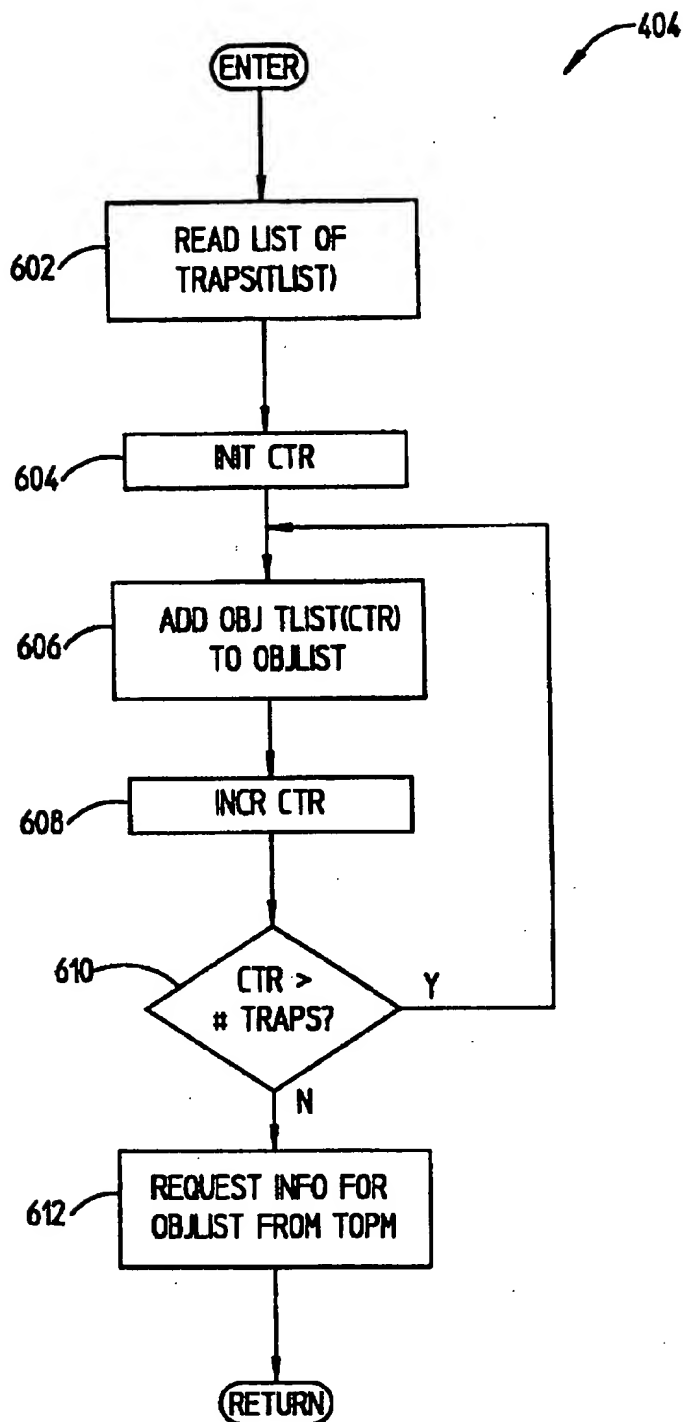


FIG 4

**FIG 5**

**FIG 6**

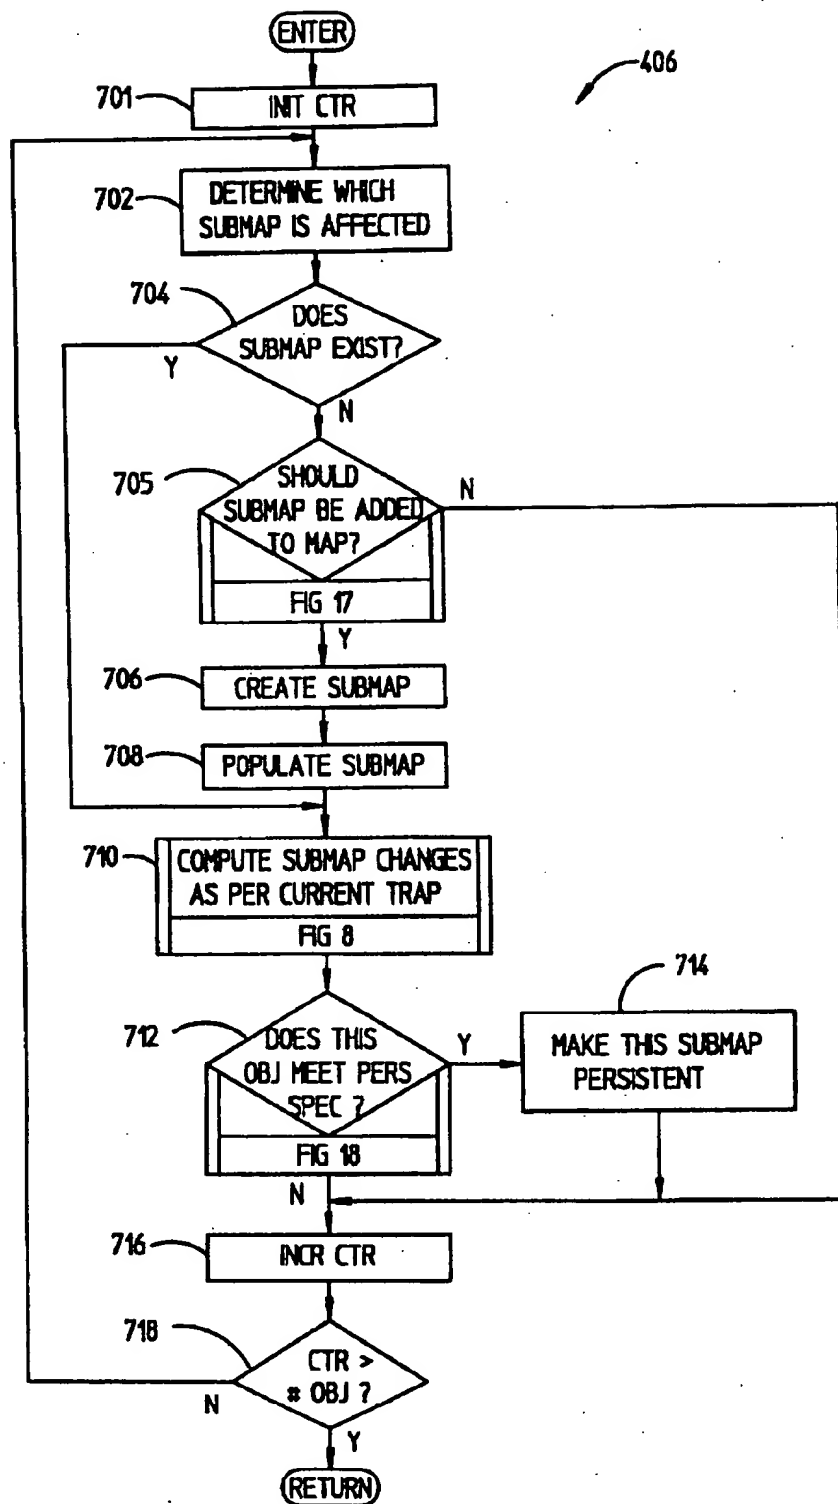
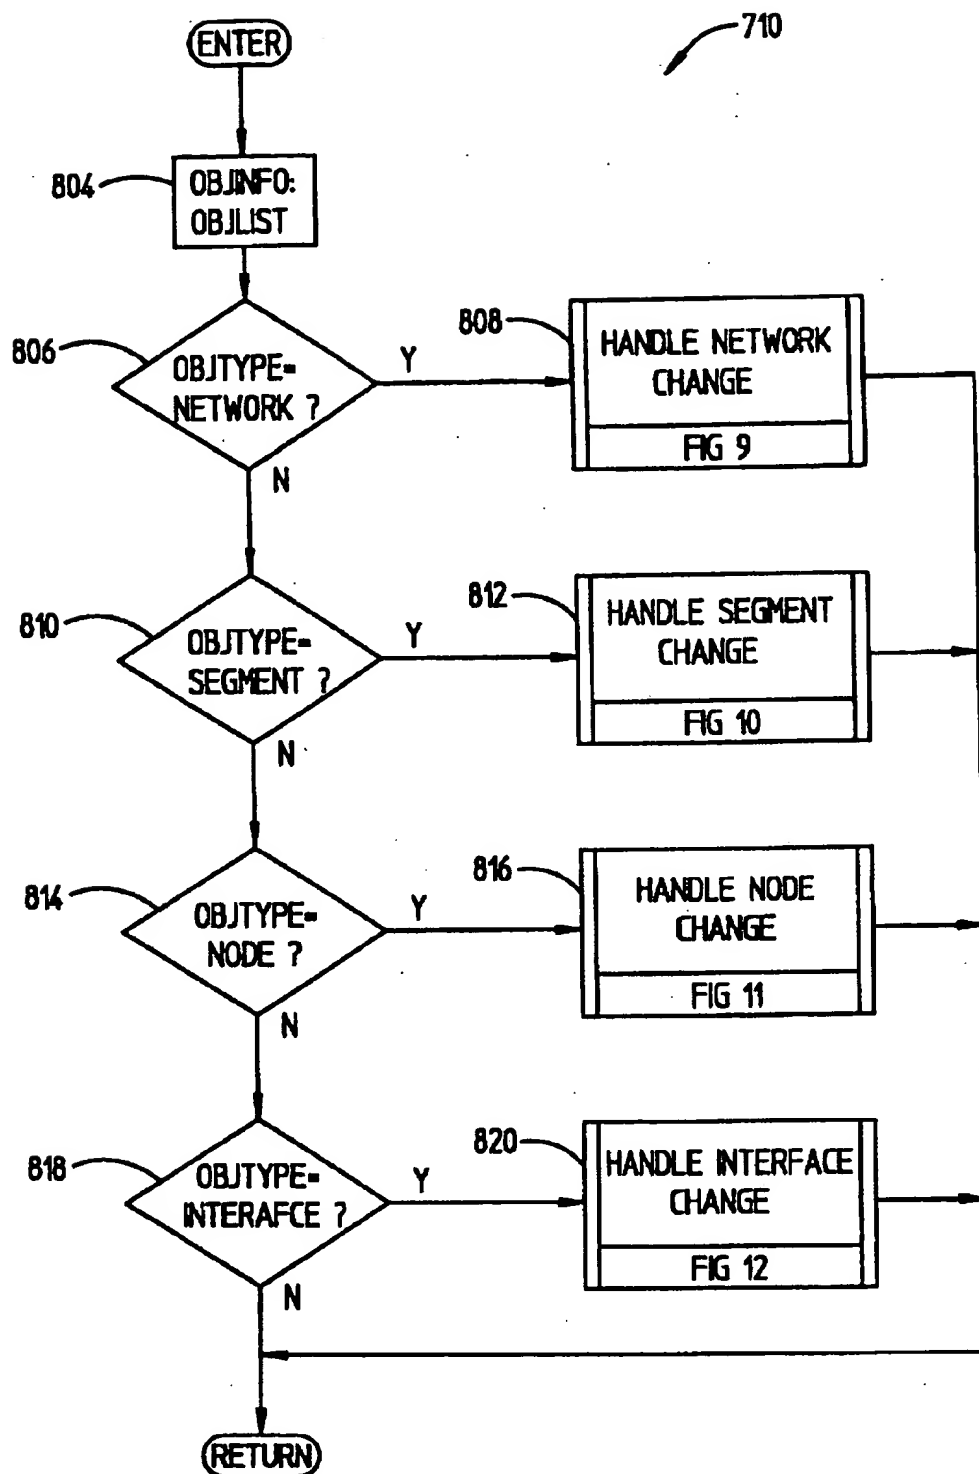
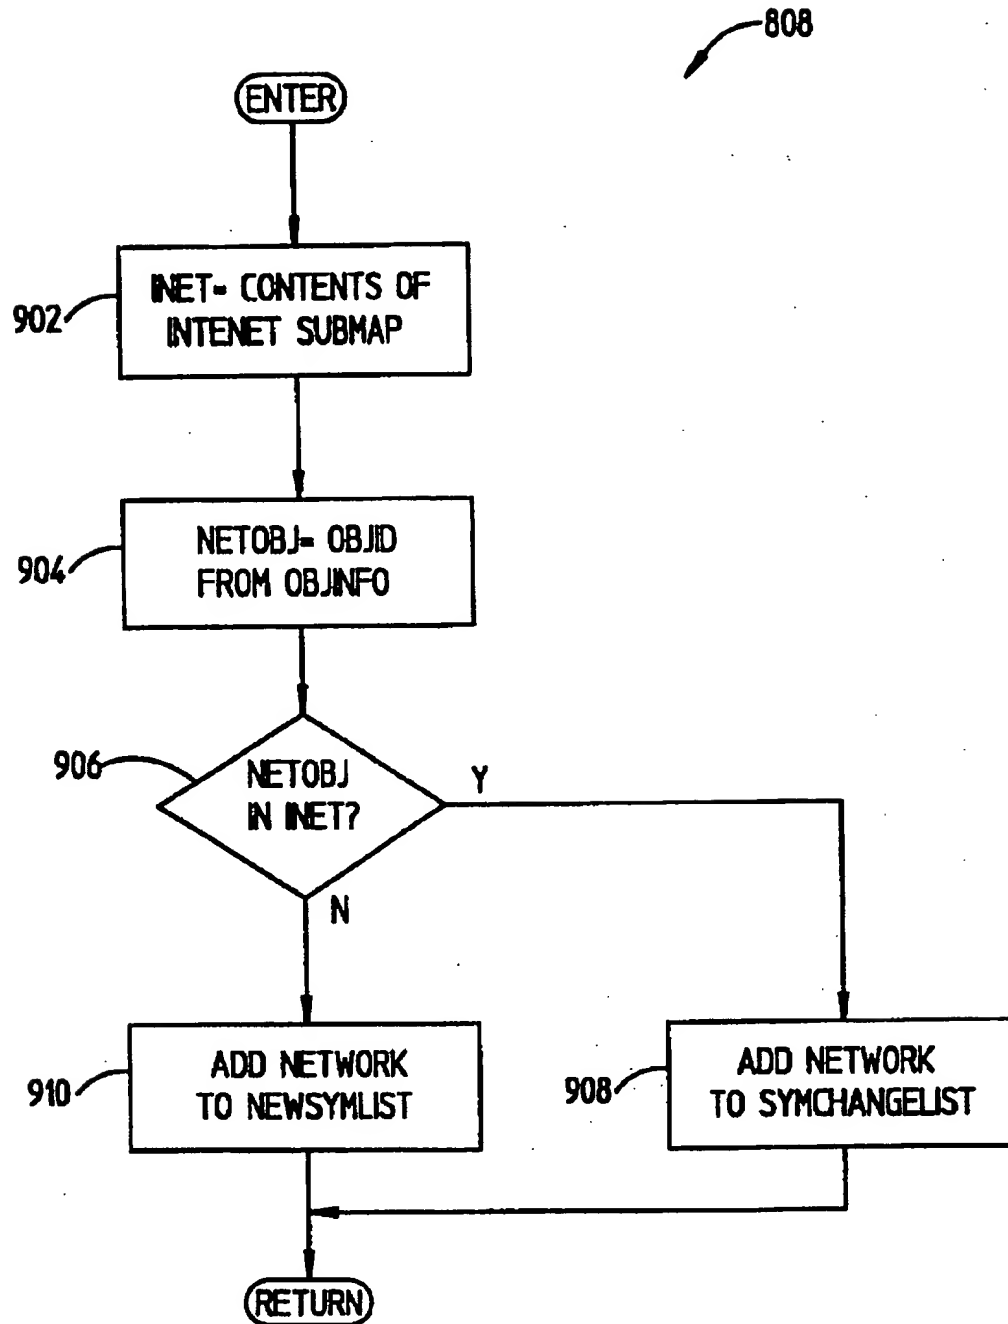


FIG 7



**FIG 8**

**FIG 9**

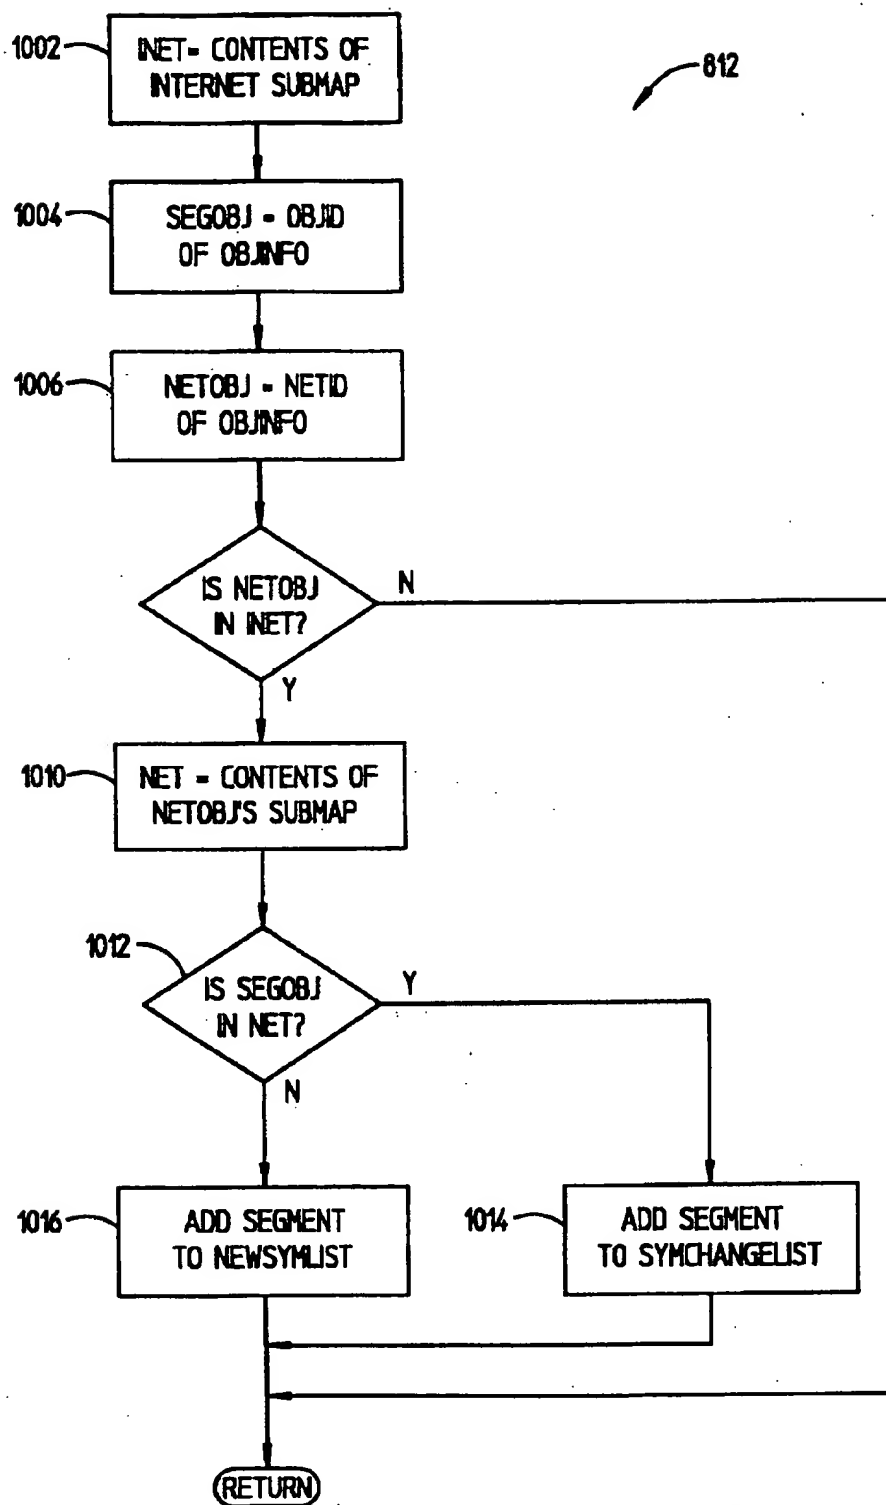


FIG 10




FIG 11

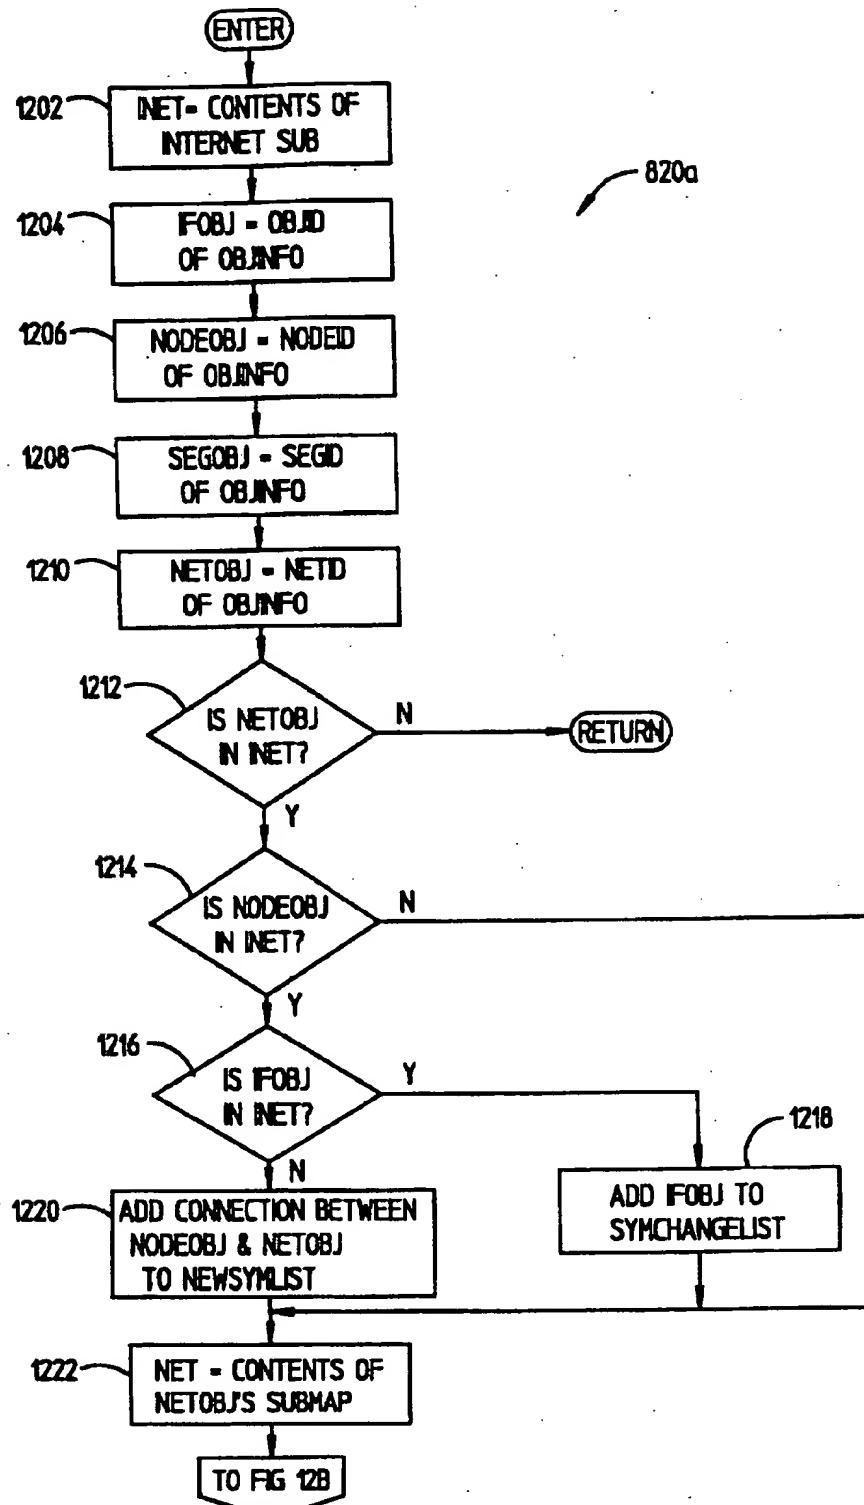


FIG 12A

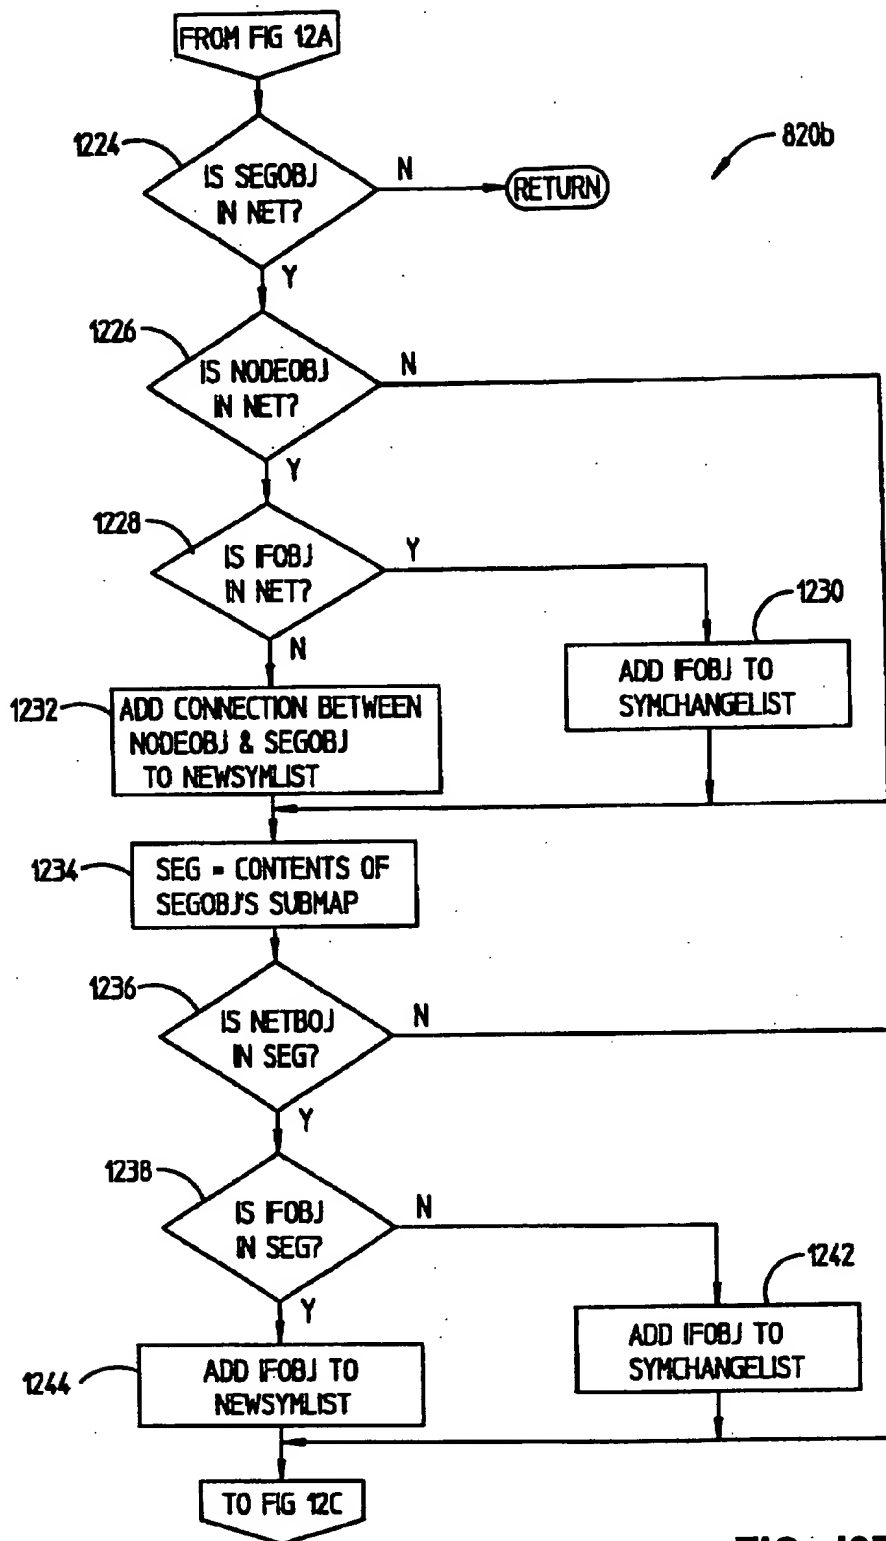
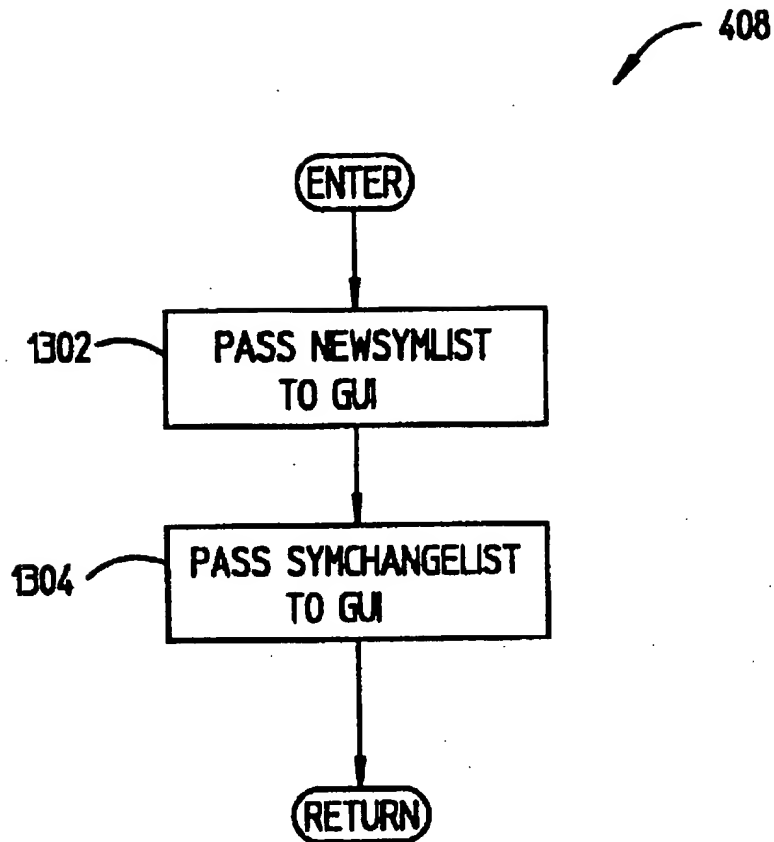
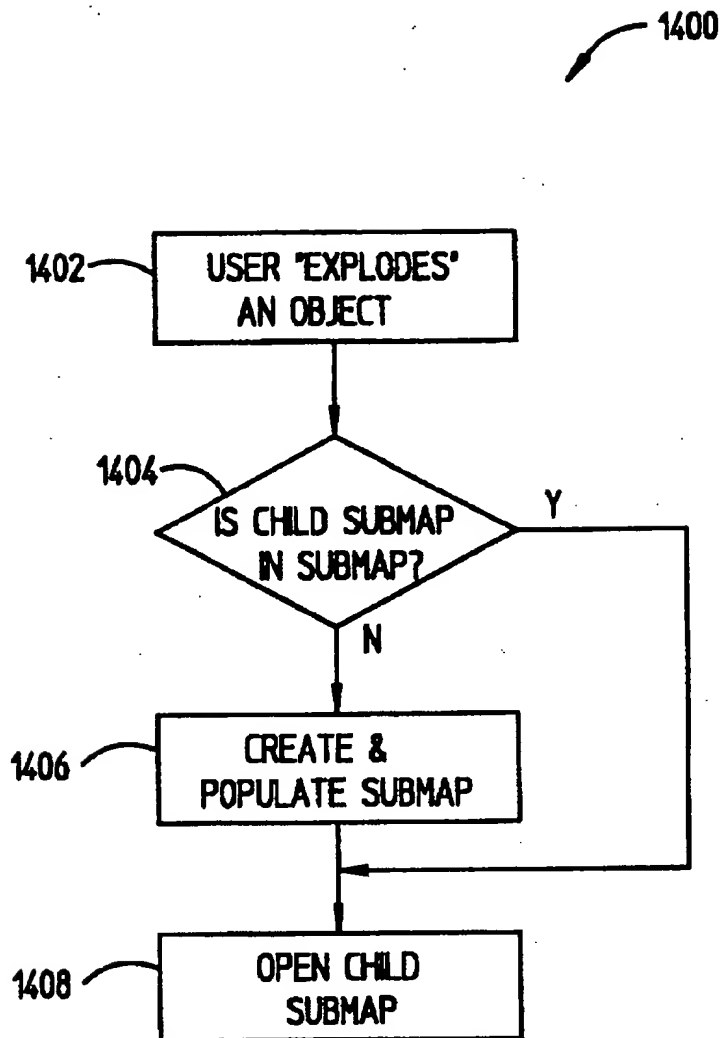


FIG 12B

**FIG 13**

**FIG 14**




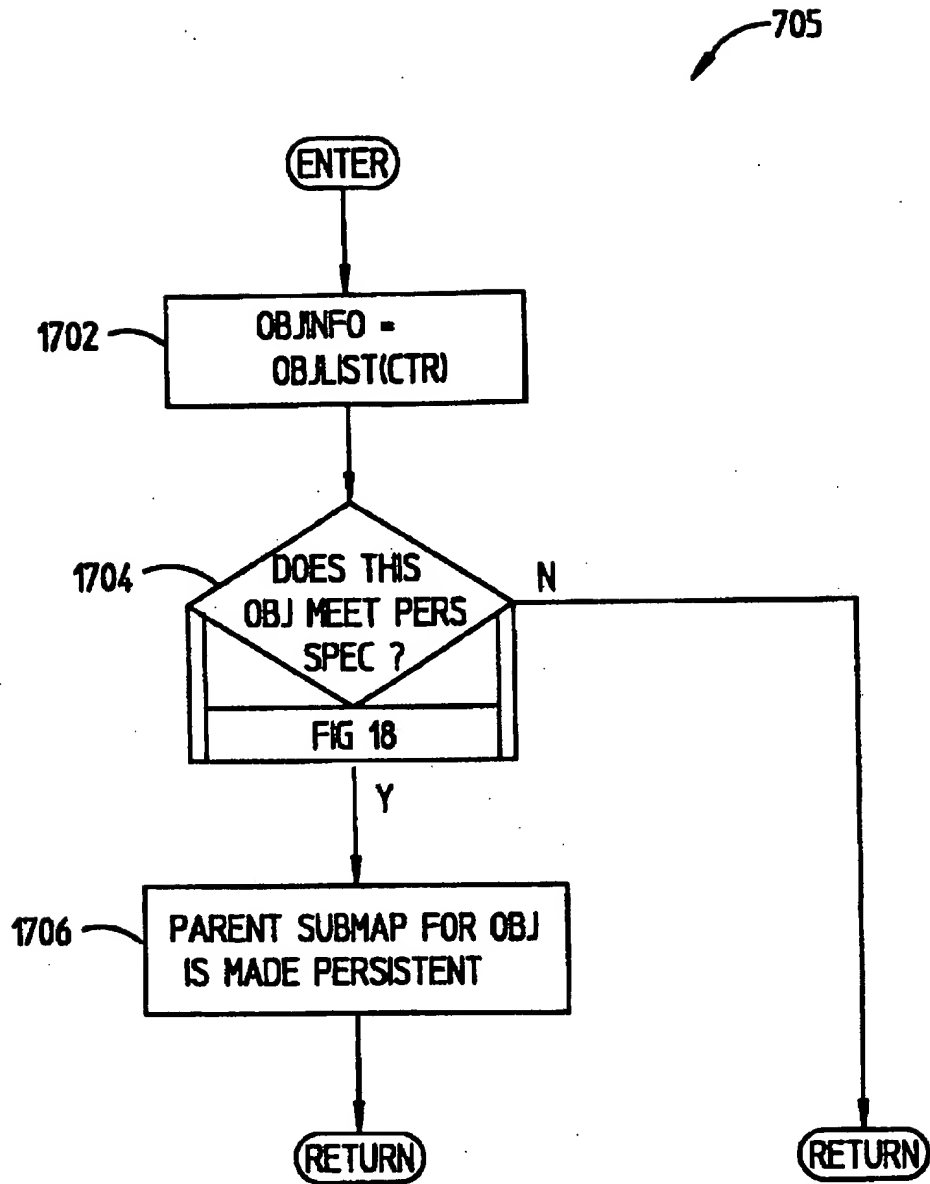


FIG 15

1600



FIG 16

**FIG 17**

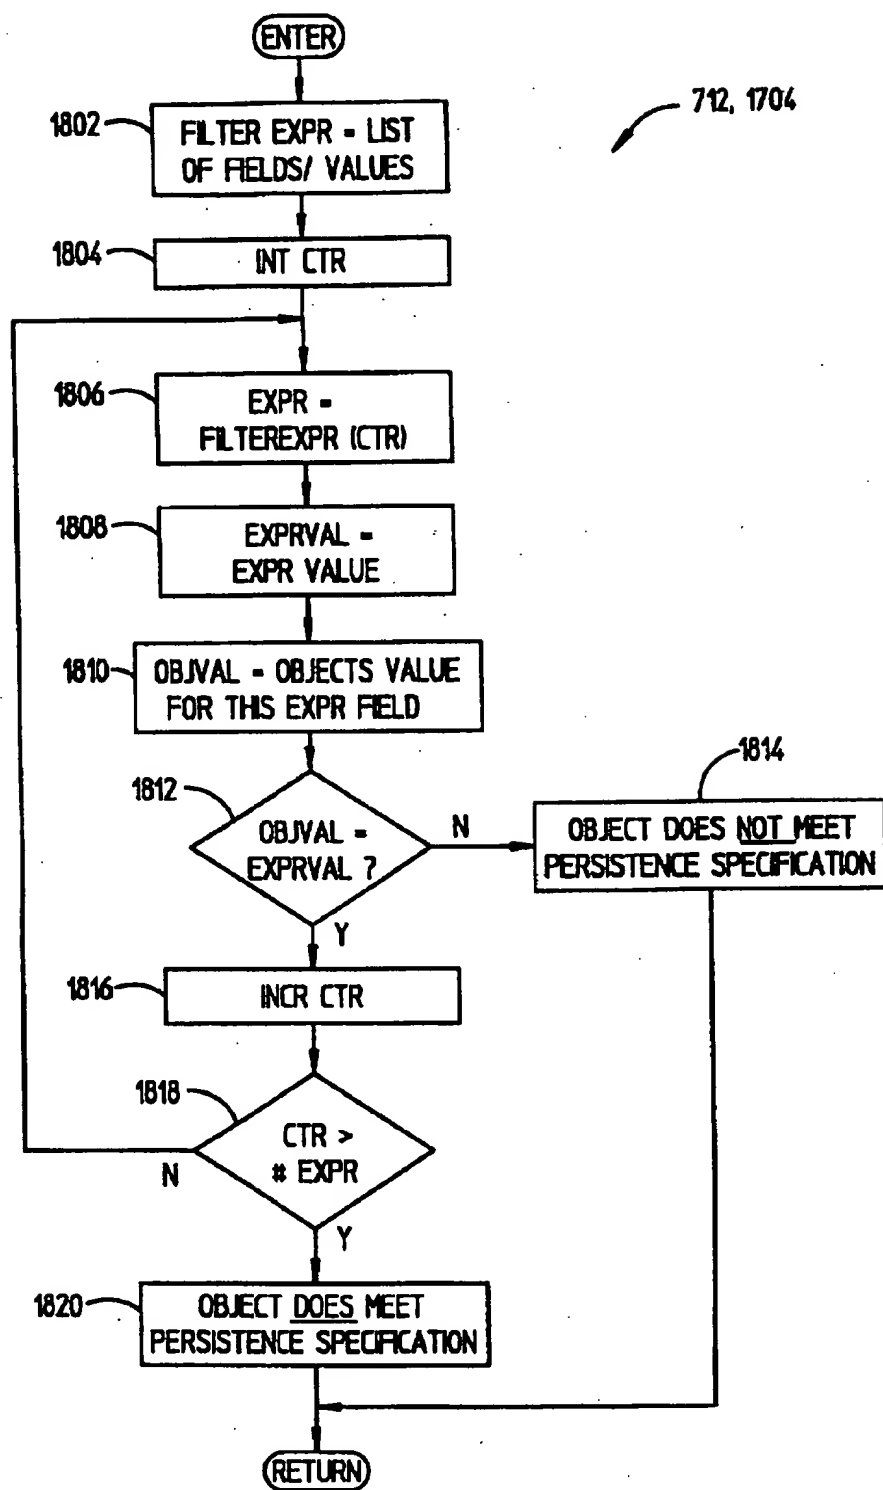


FIG 18

**PERSISTENCE SPECIFICATION SYSTEM  
AND METHOD FOR PRODUCING  
PERSISTENT AND TRANSIENT SUBMAPS IN  
A MANAGEMENT STATION FOR A DATA  
COMMUNICATION NETWORK**

**FIELD OF THE INVENTION**

The present invention relates generally to data communication networks and, more particularly, to a persistence specification system and method for permitting high performance generation of on-demand submaps in a management station for a data communication network.

**BACKGROUND OF THE INVENTION**

A data communications network generally includes a group of devices, for instance, computers, repeaters, bridges, routers, etc., situated at network nodes and a collection of communication channels for interconnecting the various nodes. Hardware and software associated with the network and particularly the devices permit the devices to exchange data electronically via the communication channels.

The size of networks varies. A local area network (LAN) is a network of devices in close proximity, typically less than one mile, and usually connected by a single cable, for instance, a coaxial cable. A wide area network (WAN) is a network of devices which are separated by longer distances, often connected by, for example, telephone lines or satellite links. In fact, some WANs span the U.S. as well as the world. Furthermore, many of these networks are widely available for use by the public, including commonly universities and commercial industries.

A very popular industry standard protocol for data communication along the networks is the Internet Protocol (IP). This protocol was originally developed by the U.S. government's Department of Defense, and has been dedicated for public use by the U.S. government. In time, the Transmission Control Protocol (TCP) and the Unreliable Datagram Protocol (UDP) were developed for use with the IP. The former protocol (TCP/IP) is a protocol which guarantees transfer of data without errors, as it implements certain check functionality, and the latter protocol (UDP/IP) is a protocol which does not guarantee transfer of data, but requires much less overhead than the TCP/IP platform. Furthermore, in order to keep track of and manage the various devices situated on a network, the Simple Network Management Protocol (SNMP) was eventually developed for use with a UDP/IP platform. The use of the foregoing protocols has become extensive in the industry, and numerous vendors now manufacture many types of network devices which can employ these protocols.

Many management software packages ("management platforms") are presently available for implementing "management stations" on a network, and particularly, in connection with the SNMP. Examples of commercially available SNMP management software packages include OpenView from the Hewlett-Packard Company (the assignee herein), NetView/6000 from IBM Corp., Spectrum from Cabletron Systems, Inc., NetLabs/Manager from NetLabs, Inc., and SunNet Manager from SunConnect, Inc. The nodes on a network and their interconnections, oftentimes referred to as the network "topology," are best displayed in a graphical format, and most, if not all, of the available management software packages provide for this feature. Typically, with these packages, a network can be viewed from different vantage points, depending on the scope of the view desired. For example, one view of the

network could be a very wide encompassing view of all nodes on the entire network. A second view could be a view of those portions of a network within a local range, for example, within a particular site or building. A third view of a network, often called a segment, could be a view of the nodes attached to a particular LAN cable.

Hewlett-Packard's very successful OpenView has been the subject of several patents, including for instance, U.S. Pat. No. 5,185,860 issued to J. C. Wu on Feb. 9, 1993, and U.S. Pat. No. 5,276,789 issued to Besaw et al on Jan. 4, 1994, the disclosures of which are incorporated herein by reference. U.S. Pat. No. 5,185,860 describes an automatic discovery system for a management station for determining the network devices and interconnections of a network, or the topology. U.S. Pat. No. 5,276,789 describes a graphic display system for a management station for graphically displaying the topology of a network and provides for various views (including, internet, segment, and node views) that can be requested by a user.

Although the presently available SNMP management stations are meritorious to an extent, the art of SNMP management stations is still in a state of infancy, and the performance of these management stations can still be enhanced and optimized. A specific area where optimization is envisioned involves the integrating applications which are associated with the management stations and which can provide additional information about network devices, for example, routers. The additional information could include, for example, device configuration information, device status, performance parameters, additional topology information, etc. However, in prior art management stations, this additional information is oftentimes not made available to the user. One reason that this additional information is not made available is that these systems suffer from an inadequate interface between the integrating applications and the display map generator within these systems.

More specifically, many management stations have on-demand submap capabilities. One such example is Hewlett-Packard's OpenView. In on-demand submap systems, a submap corresponds with each view of the network to be displayed. The system map is the collection of all submaps. In these on-demand submap systems, and particularly, the OpenView system, the user specifies which submaps the user wishes to have available, and hence, specifies the submaps which are resident within the map. Moreover, the user can also open, or "explode," a submap during operation, even though it is not specified as resident in the map. In this case, the submap is generated immediately from topology data when the user prompts the management station to open the submap, hence the name on-demand. However, the foregoing design is problematic. First, if a submap is not resident in the map, and the user opens the submap, then an integrating application cannot adequately determine what additional information to provide to the management station, due to time constraints in opening the desired submap, and the user is therefore not advised of this additional information. In other words, the integrating applications need time to analyze the topology data and time to generate the additional information for display to the user.

Another problem with current designs is that for nonresident submaps, integrating applications are unable to alert the user of critical information until the user explodes the submap. It is desirable to alert the user as soon as an anomaly is detected.

Furthermore, in prior art management stations, once a submap has been created, it remains in the system's map

from then on, even if the integrating applications no longer need the submap. This predicament undesirably uses valuable memory space. This situation also needlessly expends processor time, as some events, or changes in network topology, will unnecessarily be processed for updating some of the unnecessary resident submaps. Thus, a need exists in the industry for a system and method for better interfacing integrating applications to an on-demand submap system, for minimizing memory requirements thereof, and for optimizing performance (including speed) thereof.

### SUMMARY OF THE INVENTION

An object of the present invention is to overcome the deficiencies and inadequacies of the prior art as noted above and as generally known in the industry.

Another object of the present invention is to provide a system and method for enhancing the performance of an on-demand submap process in a management station.

Another object of the present invention is to provide optimized interfacing of integrating applications with submap information.

Another object of the present invention is to provide a system and method for minimizing memory requirements for an on-demand submap process in a management station.

Another object of the present invention is to provide a system and method for minimizing requisite processing time for an on-demand submap process in a management station.

Briefly described, the present invention is a persistence specification system and method for a management station or other device of a network for enhancing an interface between an integrating application and a graphical user interface so that more information concerning a network is provided to a user, while memory requirements and processing time are minimized. The system comprises a processor which executes the various software elements of the system, a discovery mechanism for determining the network topology data, and a layout mechanism for converting the network topology data to map data and for driving a display with the map data.

More specifically, the discovery mechanism has a network monitor configured to monitor and determine the topology of the network, a topology database for storing the topology data, and a topology manager for managing the topology database. The network monitor is also configured to generate topology change events when a change in network topology occurs and also to receive such events from other devices connected to the network.

The layout mechanism is in communication with the discovery mechanism and one or more integrating applications. The layout mechanism has a topology-to-map translator, a graphical user interface (GUI), and a map database. The translator is configured to convert the topology data to the map data, to receive the events, and to change the map data based upon the events. The GUI is configured to receive the map data from the translator, to manage the map data in the map database, and to drive the display based upon the map data. The GUI also provides information to and receives information from the integrating applications concerning the map data.

In the preferred embodiment, the translator is configured to generate a set of hierarchically arranged submaps corresponding with various views of the network. A map is defined herein as the collection of all submaps. The hierarchically arranged submaps include an internet submap having at least one network object, at least one network submap

associated with the at least one network object and having at least one segment object, at least one segment submap associated with the at least one segment object and having at least one node object, and at least one node submap associated with the at least one node object and having at least one interface object.

In accordance with a significant feature of the present invention, the translator has a persistence specification mechanism. The persistence specification mechanism and associated methodology specify when an object to be displayed is "persistent" and when an object to be displayed is "transient" based upon inputs corresponding to the object. The inputs are generated by the user and/or by the integrating applications which are connected to the discovery/layout software. The translator defines submaps as persistent when the submap has a persistent object and as transient when the submap is without persistent objects. Persistent submaps are generated and maintained continuously in the translator, so that integrating applications effectively have continuous access and can use the persistent object information to produce and provide additional information for the user. Transient submaps are only generated for a temporary time period when a user prompts the system, and the translator is configured to delete (not store) the transient submaps after the user is through viewing the transient submaps. In conclusion, the persistence specification system enhances the interface between the integrating applications and the graphical user interface so that more information concerning a network is provided to the user. The system further minimizes memory requirements and processing time by removing unnecessary submaps which could be the subject of change based upon incoming events.

In addition to achieving all of the aforementioned objects, the present invention has numerous other advantages, a few of which are delineated hereafter.

An advantage of the present invention is that it is simple in design and easy to implement on a mass commercial scale.

Another advantage of the present invention is that it is efficient as well as reliable in operation.

Another advantage of the present invention is that the principles of the persistence specification system and method can be applied not only to management stations, but also to other devices, including devices which do not practice the SNMP.

Other objects, features, and advantages of the present invention will become apparent to one of skill in the art upon examination of the following drawings and detailed description. All such additional objects, features, and advantages are intended to be included herein within the scope of the present invention.

### BRIEF DESCRIPTION OF THE DRAWINGS

The present invention can be better understood with reference to the following drawings taken in the context of the text.

FIG. 1 is a block diagram of a management station having discovery/layout software which employs the persistence specification system and method of the present invention;

FIG. 2 is a schematic diagram illustrating a display map, which comprises a collection of submaps, any of which can be displayed on the display of the management station by the discovery/layout software of FIG. 1;

FIG. 3 is a block diagram of the discovery/layout software of FIG. 1 and an integrating application interfaced therewith in accordance with the present invention;

FIG. 4 is a flow chart illustrating the architecture of the topology-to-map translator of FIG. 3;

FIG. 5 is a flow chart illustrating the architecture of a read batch block of FIG. 4;

FIG. 6 is a flow chart illustrating the architecture of a retrieve objects block of FIG. 4;

FIG. 7 is a flow chart illustrating the architecture of the compute map changes block of FIG. 4;

FIG. 8 is a flow chart illustrating the architecture of a compute submap changes block of FIG. 7;

FIG. 9 is a flow chart illustrating the architecture of a network change block of FIG. 8;

FIG. 10 is a flow chart illustrating the architecture of a segment change block of FIG. 8;

FIG. 11 is a flow chart illustrating the architecture of a node change block of FIG. 8;

FIGS. 12A and 12B is a flow chart illustrating the architecture of an interface change block of FIG. 8;

FIG. 13 is a flow chart illustrating the architecture of an update map block of FIG. 4;

FIG. 14 is a flow chart illustrating the architecture of an on-demand submap block within the graphical user interface (GUI) of FIG. 2;

FIG. 15 is a state diagram corresponding to a submap of FIG. 2 in accordance with the novel persistence specification system and method of FIG. 1;

FIG. 16 is a schematic diagram of a persistence vector corresponding to a submap for implementing the persistence specification system and method of FIG. 1;

FIG. 17 is a flow chart illustrating the architecture of a submap add decisional block of FIG. 7 for implementing the persistence specification system and method of FIG. 1; and

FIG. 18 is a flow chart illustrating the architecture of a persistence specification decisional block of both FIG. 7 and FIG. 17 for implementing the persistence specification system and method of FIG. 1.

#### DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

The following description is of the best presently contemplated mode of carrying out the present invention. This description is not to be taken in a limiting sense, but is made merely for the purpose of describing the general principles of the invention. The scope of the invention should be determined by referencing the appended claims.

FIG. 1 shows a block diagram of an object-oriented management station 100 which is implemented with a general purpose computer system containing novel discovery/layout software 101, which employs a persistence specification mechanism 103 and associated methodology of the present invention. With reference to FIG. 1, the management station 100 contains a conventional processor 102. The processor 102 communicates to other elements within the management station 100 over a system bus 104. An input device 106, for example, a keyboard or mouse, is used to input data from a user of the system 100, and a display 108 is used to output data to the user. A network interface 112 is used to interface the management station 100 to a network 118 in order to allow the management station 100 to act as a node on a network 118. A disk 114 may be used to store the software of the discovery/layout software 101 of the present invention, as well as to store the databases (topology and map) generated by the discovery/layout software 101. A printer 116 can be used to provide a hard copy output of the

nodes of the network 118 discovered by the discovery/layout software 101. A main memory 110 within the management station 100 contains the discovery/layout software 101. The discovery/layout software 101 communicates with a conventional operating system 122 and conventional network software 124 to discover the nodes on the network 118. The network software 124 serves as the intelligence, including validation, for the data communication protocols. As shown in FIG. 1, in the preferred embodiment, the network software implements the IP, the TCP and UDP over the IP, and the SNMP over the UDP. All of the foregoing protocols are well known in the art.

The discovery/layout software 101 implements object-oriented functionality. In the context of SNMP managers, object-oriented means that most of the management system actions and processes that the user can invoke are oriented toward a class of devices rather than individually managed network nodes.

Generally, the discovery/layout software 101 of FIG. 1 is configured to discover the network topology, that is, all network nodes and node interconnections existing on the network 118, and to construct a map, comprising various submaps, any of which can be used for displaying the network topology on the display 108. FIG. 2 shows a map 200 which is generated by the discovery/layout software 101 from topology data discovered from the network 118. The discovery/layout software 101 can drive any of the various submaps to the display 108 (FIG. 1) for viewing by the user.

The submaps in the map 200 of FIG. 2 are arranged in a hierarchy. A root submap 202 is defined at a root level. The root submap 202 represents the highest logical level submap in the hierarchy and shows objects 203 acting as anchor points for different submap hierarchies. Each hierarchy is a separate management domain. This could be, for instance, a network, logical grouping of nodes, or some other domain. An internet submap 204 is defined at an internet level and is generated by "exploding" an object 203 within the root submap 202. "Exploding" in the context of this document means that the user prompts the management station 100 with the input device 106 to break down and provide more data pertaining to the object 203 at issue. Further, the internet submap 204 illustrates objects 203 in the form of networks and routers. Any one of a number of network submaps 206 can be exploded from the internet submap 204. Each network submap 206 shows objects 203 in the form of segments and connectors. Any one of a number of segment submaps 208 can be exploded from an object 203 within a network submap 206. Each segment submap 208 shows objects in the form of network nodes. Finally, any one of a number of node submaps 210 can be exploded from an object 203 within a segment submap 208. Each node submap 210 shows objects 203 in the form of interfaces within that node.

In the preferred embodiment, although not necessary to practice the present invention, the discovery/layout software 101 implements on-demand submaps in order to save memory and processing time. The concept of on-demand submaps is to only place those submaps in the map 200 of FIG. 2 which the user wants to see. The net result is that only a portion of the submap hierarchy is in the map 200 at a given time. In FIG. 2, submaps (nonresident) which are not present, but would be created upon prompting by the user, are indicated by hatching. The resident submap subset of the hierarchy will change over time as the user traverses the submap hierarchy and causes nonresident submaps to be created.

A high level block diagram of the discovery/layout software 101 (FIG. 1) is set forth in FIG. 3. With the exception

of the persistence specification mechanism 103, the architecture of the discovery/layout software 101 in FIG. 3 is essentially the same as or similar to the architecture of Hewlett-Packard's well known and commercially available management software package called OpenView. As shown in FIG. 3, at a general architecture level, the discovery/layout software 101 comprises a discovery mechanism 302 for discovering nodes and interconnections of the network 118 and a layout mechanism 304 for receiving topology data from the discovery mechanism 302 and for generating the map 200 (FIG. 2) for driving the display 108. Moreover, one or more integrating applications 332 may communicate display and map information with the layout mechanism 304.

The discovery mechanism 302 has a network monitor 306 connected to the network 118 as indicated by connections 308a, 308b, a topology manager 310 connected to the network monitor 306 as indicated by arrows 312a, 312b, and a topology database in communication with the topology manager 310 as indicated by arrow 316.

The network monitor 306 transmits and receives data packets to and from the network 118. The network monitor 306 discovers and monitors network topology, as indicated by arrows 308a, 308b. When network topology changes on the network, the network monitor 306 generates events, or traps (SNMP vernacular), which include an object identifier and object change information. The network monitor 306 can also receive events from other devices, such as a router, in the network 118. The network monitor 306 interacts with the network 118 by way of the network software 124 (FIG. 1), which essentially comprises protocol stacks, corresponding to IP, TCP, UDP, and SNMP in the preferred embodiment, and which generally implements these protocols and performs validation functions. Furthermore, the network monitor 306 populates the topology database 314 by way of the topology manager 310 and notifies the topology manager 310 of events (topology changes). Finally, it should be noted that U.S. Pat. No. 5,185,860 to Wu, which is incorporated herein by reference, describes a node discovery system which could be employed to implement the network monitor 306 herein.

The topology manager 310 manages the topology database 314, as indicated by bidirectional arrow 316. The topology manager 310 prompts the network monitor 306 to update topology data related to particular events, as indicated by arrow 312a, and receives topology updates, as indicated by arrow 312b.

The topology database 314 stores topology data based upon objects, which are used to partition the network for logical reasons. Objects include, for example but not limited to, a network, a segment, a computer, a router, a repeater, a bridge, etc. Moreover, the topology data stored with respect to the objects includes, for example but not limited to, an interface or device address, an interface or device type, an interface or device manufacturer, and whether an interface or device supports the SNMP.

The layout mechanism 304 has a topology-to-map translator 318 in communication with the topology manager 310 as indicated by arrows 320a, 320b, a graphical user interface (GUI) 322 in communication with the topology-to-map translator 318 as indicated by arrows 324a, 324b, and a map database 326 in communication with the GUI 322 as indicated by bidirectional arrow 328. The integrating application 332 communicates information with the GUI 322, as indicated by arrows 333a, 333b.

It should be noted that the network monitor 306, the topology manager 310, the translator 318, and the GUI 322

take turns utilizing the combination of the operating system 122 (FIG. 1) and the processor 102 (FIG. 1) in order to accomplish their respective functions. A "context switch" as used herein refers to a change in control of the system 122 and/or processor 102 by the foregoing software elements.

The translator 318 converts topology data from the topology database 314 to map data and constructs the various submaps 202-210 in the map 200 of FIG. 2. The translator 318 can forward a request to the topology manager 310, as indicated by arrow 320a, in order to obtain topology data regarding particular objects. Moreover, in addition to forwarding topology data to the translator 318 upon request, the topology manager 310 advises the translator 318, as indicated by the arrow 320b, when topology data has changed based upon an event so that the translator 318 can make any appropriate changes in the submaps.

The GUI 322 manages the map database 326, as indicated by the bidirectional arrow 328, and manages the display 108 and input device 106, as indicated by the arrows 330a, 330b. The GUI 322 receives map updates from the translator 318, as indicated by arrow 324b, and submits user-triggered events to the translator 318, as indicated by arrow 324a. A user-triggered event includes a prompt 330a from a user to explode an object, as described relative to FIG. 2. Finally, it should be noted that U.S. Pat. No. 5,276,789 to Besaw et al., which is incorporated herein by reference, describes a graphical user interface which could be employed to implement the GUI 322 herein.

FIG. 4 shows a flow chart 400 indicating the architecture and functionality of the preferred embodiment of the topology-to-map translator 318 (FIG. 3). The translator employs the persistence specification mechanism 103 and associated methodology in accordance with the present invention, which minimize context switches (changes in the control of the operating system 122 and/or processor 102) and significantly enhance the speed and performance of the discovery/layout software 101.

With reference to FIG. 4, initially, events are queued and accumulated in a queue (not shown), or accumulator, associated with the topology manager 310, and await retrieval by the translator 318. The translator 318 reads a batch of events from the topology manager 310 during each access. Each event contains an object identifier and an object change. Moreover, the batch is any number of events greater than zero. In the tested embodiment, the batch was limited to a number no greater than 500 events, but other settings, either less than or greater than (perhaps significantly) this number could be utilized, depending upon the configuration of the system.

Next, as indicated in block 404, the translator 318 calls the topology manager 310 for a list of topology data regarding all objects which were identified in the events. After receiving the topology data, block 404 transfers to block 406.

At block 406, the translator 318 computes the changes to be made to the map data, particularly the map 200 (FIG. 2), based upon the topology data changes indicated in the events. Block 406 transfers to block 408.

At block 408, the translator 318 updates the map 200 (FIG. 2) by calling the GUI 322 and advising the GUI 322 of all submap changes (SYMBOLCHANGELIST and NEWSYMBOLLIST described hereinafter) pertaining to all object changes. This transaction is preferably, although not necessarily, a batch transfer. During this batch transfer transaction, the translator 318 identifies each submap to be changed, each object to be changed within a submap, and the particular change to be effectuated to the object. An object

change may include, for example but not limited to, a color, position, or connection change. Block 408 transfers to block 410.

At block 410, the translator 318 determines whether there is another batch of events to be read from the topology manager 310. If so, then block 410 transfers to block 402 and the previously described process is repeated. If not, then the software waits at block 410 for another batch of events.

Because of the preferred embodiment of the translator 318 set forth in FIG. 4, topology data pertaining to various objects is retrieved in a batch from the topology manager 310 and, furthermore, map data pertaining to various submaps is transferred in a batch from the translator 318 to the GUI 322. The foregoing implementation minimizes context switches, i.e., minimizes the number of times that control of the processor 102 (FIG. 1) and/or the operating system 122 (FIG. 1) is passed from one software module to another.

FIG. 5 shows a flow chart of the architecture and functionality for implementing a preferred embodiment of the read batch block 402 (FIG. 4). This flow chart illustrates how the translator 318 reads a batch of events from the topology manager 310. As indicated in a block 502, initially, events from the topology manager 310 which indicate changes in topology data are accumulated (queued). A counter at block 504 is used in connection with a loop in order to route each event from the topology manager 310 to the translator 318. At block 506, an event is read by the translator 318 from the manager 310. Block 506 transfers to block 508, which decodes the event. The event is decoded to identify the type of event and associated data. There are numerous types of events, and different types of events will have different types of associated data. More specifically, an event can involve, for example but not limited to, a new node or a node status change (e.g., connected/accessible or connected/unaccessible). An event has an event identifier, usually at the header, for identifying the type of event. Moreover, in the case of a new node, the event will contain an object identifier and an address. In the case of a node status change, the event will contain an object identifier, the old status, and the new status.

Block 508 transfers to block 510. At block 510, the decoded event data (i.e., a record) is added to a TLIST. At block 512, the counter is incremented so that another event is serviced. Block 512 transfers to block 514, which determines whether there are any more events to be serviced. If so, then block 514 transfers back to block 506 and the aforementioned process is repeated. If not, then block 514 returns to block 404 (FIG. 4).

FIG. 6 shows a flow chart of the architecture and functionality of a preferred embodiment for implementing the retrieve object information block 404 (FIG. 4). As indicated in FIG. 6, in this flow chart, object information (OBJINFO) is decoded from the decoded event data contained in the TLIST. At block 602, the TLIST is read. Block 602 transfers to block 604, which initiates an event counter. The counter in connection with a loop causes all of the events within the TLIST to be serviced. In the loop, at block 606, a single record is read. From the record, (a) an object identifier and (b) an object change are determined. The foregoing data is placed in an object list (OBJLIST). Next, at block 608, the counter is incremented so that another record of TLIST is serviced, if any remain. Block 608 transfers to block 610. At block 610, it is determined whether there are any events left to be serviced by comparing the record count of the record counter to the total number of records already processed. If so, then block 610 transfers back to block 606, which begins

to service another record. If not, then the block 610 transfers to block 612, which sends a request to the topology manager 310 for a batch transfer of object information pertaining to all of the objects within the batch. FIG. 7 shows a flow chart of the architecture and functionality of a preferred embodiment of the compute map changes block 406 (FIG. 4). In this flow chart, the translator determines which submaps (FIG. 2) are changed and the change to be effectuated, based upon the object identifiers and the object changes, which were previously determined based upon the events. With reference to FIG. 7, block 701 initiates an object change counter so that all object changes are considered. Block 701 transfers to block 702. Block 702 determines a submap identifier based upon which of the submaps (FIG. 2) are affected by the object change which is presently at issue. Block 702 transfers to block 704, which determines whether the affected submap exists. If the submap does exist, then the block 704 transfers to the block 710. If the submap does not exist, then the block 704 transfers to the block 705. At block 705, a determination is made as to whether the submap at issue should be created based upon the persistence specification of the present invention, i.e., whether the submap contains persistent objects. The persistence specification will be more fully described relative to FIGS. 15 through 18 hereinafter (particularly, FIG. 17 is a flow chart illustrating the specific architecture of block 705). If the submap is not to be added, then the block 705 transfers to block 716. If the submap is to be added, then block 705 transfers to block 706, which creates the affected submap in the map 200 (FIG. 2). Block 706 transfers to block 708.

At block 708, the translator 318 populates the newly created submap with map data retrieved from the topology database 326. Next, at block 710, submap changes based upon the current event, particularly the object identifier and the object change, are computed. These computations of block 710 will be described hereinafter relative to FIG. 8. Block 710 transfers to block 712, which makes a determination as to whether the object at issue meets the novel persistence specification. Again, the persistence specification will be more fully described hereinafter relative to FIGS. 15 through 18 (particularly, FIG. 18 is a flow chart illustrating the specific architecture of block 712). If the object meets the persistence specification, then block 712 transfers to block 714, which identifies the submap at issue as persistent, and then block 714 transfers to block 716. If the object does not meet the persistence specification, then block 712 transfers to block 716.

At block 716, the object change counter is incremented so that another object change is considered with respect to the submaps. Block 716 transfers to block 718, which makes a determination as to whether any object changes remain to be serviced. If so, then block 718 transfers back to block 702. If not, then the flow chart terminates after block 718.

Hence, at the conclusion of the operation of the logic in FIG. 7, a batch of submap identifiers with associated submap changes has been generated from a batch of object identifiers with associated object changes.

With reference to FIG. 8, relative to the submap change computations of block 710 (FIG. 7), block 804 retrieves data concerning a single object from OBJLIST. Block 804 transfers to block 806, which determines whether the object type is a network. If so, then block 806 transfers to block 808 (flow chart in FIG. 9), which computes the submap changes, and then block 808 transfers to block 822. If not, then the block 806 transfers to the block 810.

At block 810, a determination is made as to whether the object type is a segment. If so, then the block 810 transfers



to the block 812 (flow chart of FIG. 10), which computes the segment changes to the submaps, and then block 812 transfers to block 822. If not, then the block 810 transfers to the block 814.

At block 814, a determination is made as to whether the object type is a node. If so, then the block 814 transfers to the block 816 (flow chart of FIG. 11), which computes the node changes for the submaps, and then block 816 transfers to block 822. If not, then the block 814 transfers to the block 818.

At block 818, a determination is made as to whether the object type is an interface. If so, then the block 818 transfers to the block 820 (flow chart of FIG. 12), which computes the interface changes to the submap, and then block 820 transfers to block 822. If not, then the flow chart terminates.

FIG. 9 shows a flow chart of the architecture and functionality of a preferred embodiment for implementing the network change block 808 (FIG. 8). This flow chart computes changes to the internet submap 204 (FIG. 2), which displays the networks. Moreover, note that there is only a single submap (multiple submaps are possible) at the internet level in the preferred embodiment. With reference to FIG. 9, at block 902, a variable INET is set to assume the contents of the internet submap 204 (FIG. 2). The contents include a list of network objects and router objects and a list of connections between the network and router objects. Block 902 transfers to block 904. At block 904, a variable NETOBJ is set to assume the value of the object identifier (OBJID). The OBJID is retrieved from the OBJINFO. Block 904 transfers to block 906. At block 906, a determination is made as to whether NETOBJ is in INET, i.e., whether the object to be changed resides within the internet submap 204 (FIG. 2). If so, then the block 906 transfers to the block 908, which adds the network pertaining to the NETOBJ to a list SYMCHANGELIST. If not, then the block 906 transfers to the block 910, which adds the network pertaining to the NETOBJ to a list NEWSYMLIST. The lists SYMCHANGELIST and NEWSYMLIST are ultimately forwarded by the translator 318 to the GUI 322 during the batch transfer therebetween.

FIG. 10 shows a flow chart of the architecture and functionality of a preferred embodiment for implementing the segment change block 812 (FIG. 8). In this flow chart, segment changes are determined and computed. With reference to FIG. 10, block 1002 sets a variable INET to assume the contents of the internet submap 204 (FIG. 2). The contents include a list of network and router objects and a list of connections between the network and router objects. Block 1002 transfers to block 1004. At block 1004, a variable SEGOBJ is set to assume the current object identifier (OBJID), which is retrieved from the object information (OBJINFO). Block 1004 transfers to block 1006. At block 1006, a variable NETOBJ is set to assume the network identified (NETID), which is determined from the OBJINFO. Block 1006 transfers to block 1008. At block 1008, a determination is made as to whether NETOBJ is in INET, i.e., whether the current network is within the current internet submap 204 (FIG. 2). If not, then the flow chart of FIG. 10 terminates. If so, then the block 1008 transfers to block 1010. At block 1010, a variable NET is set to assume the contents of the network submap 206 (FIG. 2) pertaining to NETOBJ. The contents include, for example but not limited to, a list of segment and connector objects and connections between segment and connectors. Block 1010 transfers to block 1012. At block 1012, a determination is made as to whether SEGOBJ is in the NET (i.e., is the segment in the network submap?). If so, then the block 1012

transfers to the block 1014, which adds the segment pertaining to SEGOBJ to the SYMCHANGELIST. Otherwise, if not, block 1012 transfers to the block 1016, which adds the segment pertaining to SEGOBJ to NEWSYMLIST. Finally, after blocks 1014, 1016, the flow chart of FIG. 10 terminates and operation transfers back to FIG. 8.

FIG. 11 shows a flow chart of the architecture and functionality of a preferred embodiment for implementing the node change block 816 (FIG. 8). In the flow chart of FIG. 11, node changes are determined and computed by the translator 318. As shown in FIG. 11, block 1102 sets a variable INET to assume the contents of the internet submap 204 (FIG. 2). The contents include a list of network and router objects and a list of connections between the network and router objects. Block 1102 transfers to block 1104. At block 1104, a variable NODEOBJ is set to assume the object identifier (OBJID) contained in the object information (OBJINFO). Block 1104 transfers to block 1106. At block 1106, a variable SEGOBJ is set to assume the segment identifier (SEGID) contained within the OBJINFO. Block 1106 transfers to block 1108. At block 1108, a variable NETOBJ is set to assume the network identifier (NETID) contained within the OBJINFO. Block 1108 transfers to block 1110. At block 1110, a determination is made as to whether NETOBJ is in INET (i.e., is the network in the internet submap?). If not, then the flow chart terminates. If so, then the block 1110 transfers to the block 1112. At block 1112, the variable NET is set to assume the contents of the network submap 206 (FIG. 2) pertaining to NETOBJ. The contents include, for example but not limited to, a list of segments, connectors and connections between segments and connectors. Block 1112 transfers to block 1114. At block 1114, a determination is made as to whether SEGOBJ is in NET. If not, then the flow chart terminates. If so, then the block 1114 transfers to the block 1116. At block 1116, the variable SEG is set to assume the contents of the segment submap 208 (FIG. 2) pertaining to SEGOBJ. The contents include, for example but not limited to, a list of nodes and connections between the nodes and the network. Block 1116 transfers to block 1118. At block 1118, a determination is made as to whether NODEOBJ is in SEG, i.e., whether the node object is in the present segment at issue. If so, then the block 1118 transfers to the block 1120, which adds the node pertaining to NODEOBJ to SYMCHANGELIST and then the flow chart terminates. Otherwise, if not, the block 1118 transfers to the block 1122 which adds the node pertaining to NODEOBJ to NEWSYMLIST and then the flow chart terminates.

FIGS. 12A through 12C collectively show a flow chart of the architecture and functionality of the preferred embodiment for implementing the interface change block 820 (FIG. 8). In this flow chart, interface changes in the submaps are determined and computed by the translator 318 (FIG. 3). With reference to FIG. 12A, a block 1202 sets a variable INET to assume the contents of the internet submap 204 (FIG. 2) which is currently at issue. The contents include a list of nets, routers and connections objects. Block 1202 transfers to block 1204. At block 1204, a variable IFOBJ is set to assume the OBJID contained within the OBJINFO. The block 1204 transfers to the block 1206. At block 1206, the variable NODEOBJ is set to assume the NODEID contained within the OBJINFO. Block 1206 transfers to block 1208. At block 1208, the variable SEGOBJ is set to assume the SEGID contained within OBJINFO. Block 1208 transfers to block 1210. At block 1210, a variable NETOBJ is set to assume the NETID contained within OBJINFO. After block 1210, the initialization process has been completed and the block 1210 transfers to block 1212.

13

At block 1212, a determination is made as to whether NETOBJ is in INET, i.e., whether the network object is in the current internet submap 204 (FIG. 2). If not, the flow chart terminates, as shown in Fig. 12A. If so, then block 1212 transfers to block 1214. At block 1214, a determination is made as to whether NODEOBJ is in INET, i.e., whether the node object is in the internet submap 204 (FIG. 2). If not, then the block 1214 transfers to the block 1222. If so, then the block 1214 transfers to the block 1216.

At block 1216, a determination is made as to whether IFOBJ is in INET. If so, then the block 1216 transfers to the block 1218, which adds the interface pertaining to IFOBJ to the SYMCHANGELIST. If not, then the block 1216 transfers to block 1220, which adds the interface pertaining to IFOBJ (between node object and network object) to NEWSYMLIST.

At block 1222, the variable NET is set to assume the contents of the network submap 206 (FIG. 2). The contents include, for example but not limited to, segments, connectors, connections, etc. Block 1222 transfers to block 1224 of FIG. 12B.

At block 1224, a determination is made as to whether SEGOBJ is in NET, i.e., whether the segment object is within the network submap 206 (FIG. 2). If not, then the flow chart terminates. If so, then the block 1224 transfers to the block 1226.

At block 1226, a determination is made as to whether NODEOBJ is in NET, i.e., whether the node object is within the network submap 206 (FIG. 2). If not, then the flow chart transfers to block 1234. If so, then the block 1226 transfers to block 1228.

At block 1228, a determination is made as to whether IFOBJ is within NET, i.e., whether the interface object is within the network submap 206 (FIG. 2). If so, then the block 1228 transfers to the block 1230, which adds the interface pertaining to IFOBJ to SYMCHANGELIST. If not, then the block 1228 transfers to the block 1232, which adds the interface pertaining to IFOBJ (which is between a node object and a segment object) to NEWSYMLIST. The blocks 1230, 1232 transfer to the block 1234.

At block 1234, the variable SEG is set to assume the contents of the segment submap 208 (FIG. 2). The contents include, for example but not limited to, nodes and connections to network (interfaces). Block 1234 transfers to block 1236.

At block 1236, a determination is made as to whether NODEOBJ is in SEG, i.e., whether the node object is within the segment submap 208 (FIG. 2). If not, then the flow chart transfers to the block 1246 of FIG. 12B. If so, then the block 1236 transfers to the block 1238.

At block 1238, a determination is made as to whether IFOBJ is within SEG, i.e., whether the interface object is within the segment submap 208 (FIG. 2). If so, then the block 1238 transfers to the block 1242, which adds the interface pertaining to IFOBJ to SYMCHANGELIST. If not, then the block 1238 transfers to the block 1244, which adds the interface pertaining to IFOBJ to NEWSYMLIST. The blocks 1242, 1244 transfer to the block 1246 of FIG. 12C.

At block 1246, the variable NODE is set to assume the contents of the node submap 210 (FIG. 2). The contents include interface objects. Block 1246 transfers to block 1248.

At block 1248, a determination is made as to whether IFOBJ is within NODE, i.e., whether the interface object is within the node submap 210 (FIG. 2). If so, then the

14

interface pertaining to IFOBJ is added to SYMCHANGELIST, as indicated at block 1250. If not, then the block 1248 transfers to the block 1252, which adds the interface pertaining to IFOBJ to NEWSYMLIST. Finally, after blocks 1250, 1252, the flow chart contained collectively in FIGS. 12A through 12C terminates.

FIG. 13 shows a flow chart of the architecture and functionality of a preferred embodiment for implementing the update map block 408 (FIG. 4). In this flow chart, a batch transfer of changes is sent by the translator 318 to the GUI 322. With reference to FIG. 13, at block 1302, the translator 318 transfers the NEWSYMLIST to the GUI 322, and in block 1304, the translator 318 transfers the SYMCHANGELIST to the GUI 322. After block 1304, the flow chart of FIG. 13 terminates and the operation passes back to block 410 (FIG. 4).

FIG. 14 illustrates an on-demand submap module contained within the GUI 322 (FIG. 3). This flow chart implements the user interface to the various submaps of the map 200 (FIG. 2). With reference to FIG. 14, at a block 1402, the GUI 322 monitors the input devices connected to the management station 100 (FIG. 1), for instance, the keyboard 106. When the user of the management station 100 prompts the management station 100 via the keyboard 106 or via some other input device to explode an object on the display 108, the block 1402 of FIG. 14 transfers to the block 1404 in order to process the user request. At block 1404, a determination is made as to whether the child submap is contained within the map 200 (FIG. 2). If so, then the block 1404 transfers to the block 1408. If not, then the block 1404 transfers to the block 1406, which creates and populates the submap. The GUI 322 populates the submap by requesting the translator 318 to create and populate a submap based on topology data retrieved from topology manager 310. Moreover, block 1406 transfers to block 1408 which opens the child submap and displays the child submap on the display 108 for the user.

#### PERSISTENCE SPECIFICATION

The concept of on-demand submaps, as illustrated in and described with respect to FIG. 2, is to only place those submaps in the map 200 (FIG. 2) which the user wants to see on the display 108 (FIG. 1). With the implementation of the persistence specification mechanism 103 (FIG. 1) herein, this concept is extended to include submaps needed by integrating applications 332 (FIG. 3), which can dynamically change during operation. The net result is that only a portion of the submap hierarchy is in the map 200 (FIG. 2) at a given time, and the information provided to the user is significantly enhanced by permitting integrating applications 332 to supplement device configuration information which has been discovered from the network 118 by the discovery mechanism 302 (FIG. 3).

The persistence specification involves defining each submap of the map 200 (FIG. 2) as either persistent or transient, as indicated in the diagram 1500 of FIG. 15. No matter what is specified by the persistence specification, the user still has access to the entire topology through the map 200. However, those submaps which are designated as persistent are placed on the map 200 immediately after this determination, and those submaps which are designated as transient are only created on demand after a user request, and then when the user leaves the transient submap, the transient submap is removed from the map 200 (FIG. 2).

The persistence specification is useful when there is an integrating application 332 (FIG. 3) which is tightly inte-

grated with the submap hierarchy. In other words, it is useful where the integrating application 332 depends on the translator 318 to place certain objects in the map 200 in order for the integrating application 332 to operate properly. For example, a company may have developed an integrating application 332 which adds a blob symbol to a router in a node submap and needs the status of the blobs to propagate up the submap hierarchy. Accordingly, the integrating application 332 would need the router specified in the persistence specification.

In order to implement the persistence specification, a persistence vector 1600 of FIG. 16 is associated with each of the submaps of the map 200 (FIG. 2) by the GUI 322 (FIG. 3). The persistence vector 1600 comprises numerous persistence bits. A persistence bit  $b_u$  corresponds to the user of the computer system 100 and is generated by the GUI 322. Moreover, there is provided a persistence bit  $b_{A1} \dots b_{An}$  corresponding with each of the integrating applications 332 (FIG. 3) that is associated with the discovery/layout software 101 (FIGS. 1, 3). In order for a submap to be considered transient, none of the persistent bits  $b_u, b_{A1}, \dots, b_{An}$  in the persistence vector 1600 must be asserted. Conversely, when any of the bits  $b_u, b_{A1}, \dots, b_{An}$  of the persistence vector 1600 is asserted, then the respective submap is considered persistent.

Referring back to the state diagram in FIG. 15, a submap transitions from transient to persistent in the following circumstances: (a) the user or an integrating application 332 makes a change (for example, adding a background graphic, moving a symbol, changing a symbol label, changing to auto or manual layout, or hiding a symbol) with respect to an object and the change does not affect anything stored in the topology database 314; or (b) an integrating application  $A_i$  asserts its corresponding persistence bit  $b_{Ai}$ . A submap changes from the persistent state to the transient state when an integrating application deasserts its corresponding persistence bit, if this results in all persistence bits of the persistence vector 1600 being deasserted. Thus, in conclusion, a user can create a persistent submap, whereas an integrating application can create either a persistent or a transient submap.

Recall that in FIG. 7, block 705 makes an inquiry as to whether the submap at issue should be added to the map 200 (FIG. 2) based upon the persistence specification. FIG. 17 is a flow chart illustrating the architecture and functionality of a preferred embodiment of block 705. As shown in FIG. 17 at block 1702, a variable OBJINFO is set to assume each of the object identifiers within the submap, and each of the object identifiers is considered via a counter. Block 1702 transfers to block 1704.

At block 1704, a determination is made as to whether the particular object at issue meets the persistence specification. If not, then block 1704 transfers to block 712 (FIG. 7). If so, then the block 1704 transfers to the block 1706, which defines the submap containing the object as a persistent submap by asserting the persistence bit (one of  $b_u, b_{A1}, \dots, b_{An}$ ) corresponding with the submap at issue. Block 1706 then transfers to block 706 (FIG. 7).

FIG. 18 shows a flow chart for determining whether an object should be classified as either persistent or transient, as is determined in block 712 (FIG. 7) and also in block 1704 (FIG. 17). Referring to FIG. 18, a block 1802 sets a variable FILTEREXPR to assume a list of fields and values, which are transferred from the translator 103 to the integrating application 322. The fields and values are essentially flags for the integrating applications 332 and identify supplement-

tal information which can be provided by the integrating applications 332. A field identifies some parameter of an object (for example but not limited to, vendor (manufacturer), device type, address, rate (e.g., packets/second), etc.). The value is simply a value for that field (for example but not limited to, CISCO, router, 15.2.112.227, 55 packets/second, etc.). So, for a vendor field, possible values are, for instance, Hewlett-Packard, CISCO, International Business Machines (IBM), etc. For a device type possible values are, for example, a router, a printer, etc.

Furthermore, block 1802 transfers to block 1804, which initiates a counter for the purpose of considering all of pairings of fields and values with respect to the object at issue. Block 1804 transfers into the loop which begins with block 1806.

At block 1806, a variable EXPR is set to assume a field and a value. Block 1806 transfers to block 1808, which sets a variable EXPRVAL to assume the value (EXPR·VALUE) within the variable EXPR. Block 1808 transfers to block 1810.

At block 1810, a variable OBJVAL is set to assume the object value pertaining to the object at issue within the particular field (EXPR·FIELD) of EXPR. Block 1810 transfers to block 1812.

At block 1812, OBJVAL is compared to EXPRVAL, i.e., the object value is compared to the integrating application value or test value. If the object value does not match the integrating application value, then the object does not meet the persistent specification, as indicated at block 1814 and the flow chart terminates. However, if the object value matches all of the integrating application values, then the object is ultimately made persistent at block 1820. However, before reaching block 1820, block 1812 transfers to block 1816, which increments the counter initiated in block 1804. Moreover, block 1816 transfers to block 1818, which determines whether all EXPRs have been considered. If some remain, then block 1818 transfers back to block 1806 and the foregoing process continues. If no more EXPRs remain, then the flow chart transfers to block 1820, which specifies the object as meeting the persistence specification and then the flow chart terminates.

When an object is definitively matched to a field/value pair and designated persistent, the translator 103 advises the GUI 322 of this fact, and then the GUI 322 informs the integrating application 332 of the existence of the persistent object. In turn, the GUI 322 provides supplemental information, as indicated by arrow 333b, pertaining to and/or based upon the persistent object to the GUI 322 for introduction into the appropriate submaps and for display.

It will be obvious to those skilled in the art that many modifications can be made to the preferred embodiment as described above without departing from the spirit and scope of the present invention. The disclosures and description are intended to be illustrative and any such modifications are intended to be included herein within the scope of the present invention, as is defined in the following claims.

Wherefore, the following is claimed:

1. A persistence specification system, comprising:

a processor;

a discovery mechanism associated with a processor, said discovery mechanism configured to generate and store topology data specifying devices and interconnections of a network; and

a layout mechanism associated with said processor and interfaced with said discovery mechanism, said layout mechanism configured to receive said topology data.

from said discovery mechanism, said layout mechanism configured to drive a display based upon said topology data, said layout mechanism comprising:

a translator configured to convert said topology data to map data, said translator having a persistence specification means, said persistence specification means for specifying based upon a persistence input when an object to be displayed is persistent and, alternatively, when said object is transient, said persistence specification means for defining said map data as persistent when said map data has a persistent object and for defining said map data as transient when said map data is without a persistent object, said translator configured to generate and continuously maintain said persistent map data and configured to generate and temporarily maintain transient data during a temporary time period corresponding to a demand by a user; and

a graphical user interface configured to receive said map data from said translator and to drive said display based upon said map data;

an integrating application in communication with said persistence specification mechanism for generating said persistence input;

said persistent specification mechanism being configured to advise said integrating application of said persistent objects; and

said integrating application being configured to provide supplemental display information to said graphical user interface based upon said persistent objects.

2. The system of claim 1, wherein said persistence input is generated by an integrating application.

3. The system of claim 1, wherein said persistence input is produced by said user.

4. The system of claim 1, wherein said translator is configured to generate a plurality of hierarchically arranged submaps from said topology data.

5. The system of claim 4, wherein said hierarchically arranged submaps include an internet submap having at least one network object, at least one network submap associated with said at least one network object and having at least one segment object, at least one segment submap associated with said at least one segment object and having at least one node object, and at least one node submap associated with said at least one node object and having at least one interface object.

6. A persistence specification system for enhancing intercommunication between an integrating application and a graphical user interface so that more information concerning a network is provided to a user, while memory requirements and processing time are minimized, comprising:

a topology database for storing a plurality of submaps of map data for a graphical user interface, said submaps for driving a display;

a translator configured to convert said topology data from said topology database to map data for said map database, said translator configured to generate submaps from said map data for said graphical user interface for driving said display;

a persistence specification mechanism associated with said translator, said persistence specification mechanism configured to evaluate objects within said map data and configured to specify each object as persistent and, alternatively, as transient, based upon said object

evaluation an persistence input pertaining to said objects, said persistence mechanism configured to evaluate each said submap and configured to specify a submap as persistent and, alternatively, as transient, based upon said submap evaluation, said submap being specified as persistent when said submap comprise a persistent object, said submap being specified as transient when said submap comprises no persistent object; and

said translator configured to generate and continuously maintain said persistent submaps within said map database, said translator configured to generate and temporarily maintain said transient submaps upon receiving a user prompt for a temporary time period corresponding with said user prompt;

said integrating application in communication with said persistence specification mechanism for generating said persistence input:

said persistence specification mechanism being configured to advise said integrating application of said persistent objects; and

said integrating application being configured to provide supplemental display information to said graphical user interface based upon said persistent objects.

7. The system of claim 6, wherein said persistence input is generated by an integrating application.

8. The system of claim 6, wherein said persistence input is produced by said user.

9. The system of claim 6, wherein said persistence input defines a vendor.

10. The system of claim 6, wherein said persistence input defines a device type.

11. The system of claim 6, wherein said translator is configured to generate a plurality of hierarchically arranged submaps from said topology data.

12. The system of claim 11, wherein said hierarchically arranged submaps include an internet submap having at least one network object, at least one network submap associated with said at least one network object and having at least one segment object, at least one segment submap associated with said at least one segment object and having at least one node object, and at least one node submap associated with said at least one node object and having at least one interface object.

13. A computer readable medium having a program for enhancing intercommunication between an integrating application and a graphical user interface so that more information concerning a network is provided to a user, comprising:

a topology database for storing topology data pertaining to devices and device interconnections of said network;

a map database for storing a plurality of submaps of map data for a graphical user interface, said submaps for driving a display;

a translator configured to convert said topology data from said topology database to map data for said map database, said translator configured to generate submaps from said map data for said graphical user interface for driving said display;

a persistence specification mechanism associated with said translator, said persistence specification mechanism configured to evaluate objects within said map data and configured to specify each object as persistent

19

and, alternatively, as transient, based upon said object evaluation and a persistence input pertaining to said objects, said persistence mechanism configured to evaluate each said submap and configured to specify a submap as persistent and, alternatively, as transient, based upon said submap evaluation, said submap being specified as persistent when said submap comprises a persistent object, said submap being specified as transient when said submap comprises no persistent object; and

10 said translator configured to generate and continuously maintain said persistent submaps within said map

20

database, said translator configured to generate and temporarily maintain said transient submaps upon receiving a user prompt for a temporary time period corresponding with said user prompt;

5 said integrating application in communication with said persistence specification mechanism for generating said persistence input;

10 said persistent specification mechanism being configured to provide supplemental information to said graphical user interface based upon said persistent objects.

\* \* \* \* \*

UNITED STATES PATENT AND TRADEMARK OFFICE  
**CERTIFICATE OF CORRECTION**

PATENT NO. : 5,689,645

DATED : November 18, 1997

INVENTOR(S) : Schettler, et al.

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Column 18, line 1, delete "*an*" and insert therefor -- *and a* --;

Column 18, line 6, delete "*comprise*" and insert therefor -- *comprises* --.

Signed and Sealed this

Thirtieth Day of January, 2001

Attest:



Q. TODD DICKINSON

Attesting Officer

Director of Patents and Trademarks